

# Energy-Aware Algorithm Selection in Particle Simulations A Case Study with AutoPas

**Manish Kumar Mishra**, Samuel J. Newcome,  
Markus Muehlhaeusser, Hans-Joachim Bungartz

Technical University of Munich  
School of Computation, Information and Technology  
Chair of Scientific Computing in Computer Science

Hamburg, 26.06.2026



Federal Ministry  
of Research, Technology  
and Space

**WindHPC**



Funded by  
the European Union




CONNECTING THE DOTS




**ISC**


High Performance




# Motivation




# Motivation




Colagrossi et al., Journal of Computational Physics 191 (2003)




# Motivation




Colagrossi et al., Journal of Computational Physics 191 (2003)




Martin et al., Ultrasound-induced granular flows: a two-time scale model, J. Fluid Mech. (2025)




# Motivation




Colagrossi et al., Journal of Computational Physics 191 (2003)



Martin et al., Ultrasound-induced granular flows: a two-time scale model, J. Fluid Mech. (2025)



Wei Wang et al., Protein folding simulations: From coarse-grained model to all-atom model, IUPMB. (2009)



# Particle simulations need HPC

# But energy consumption is a growing concern!

# Energy consumption of supercomputers

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	<b>LineShine</b> - LingKun, LX2 304C 1.55GHz, LingQi, Kylin OS, Shenzhen Cloud Computing Center Co., Ltd. National Supercomputing Centre in Shenzhen (NSCS) China	13,789,440	2,198.40	2,735.82	42,220
2	<b>El Capitan</b> - HPE Cray EX255a, AMD 4th Gen EPYC 24C 1.8GHz, AMD Instinct MI300A, Slingshot-11, TOSS, HPE DOE/NNSA/LLNL United States	11,340,000	1,809.00	2,821.10	29,685

Source: Top500 list, Released June 2026

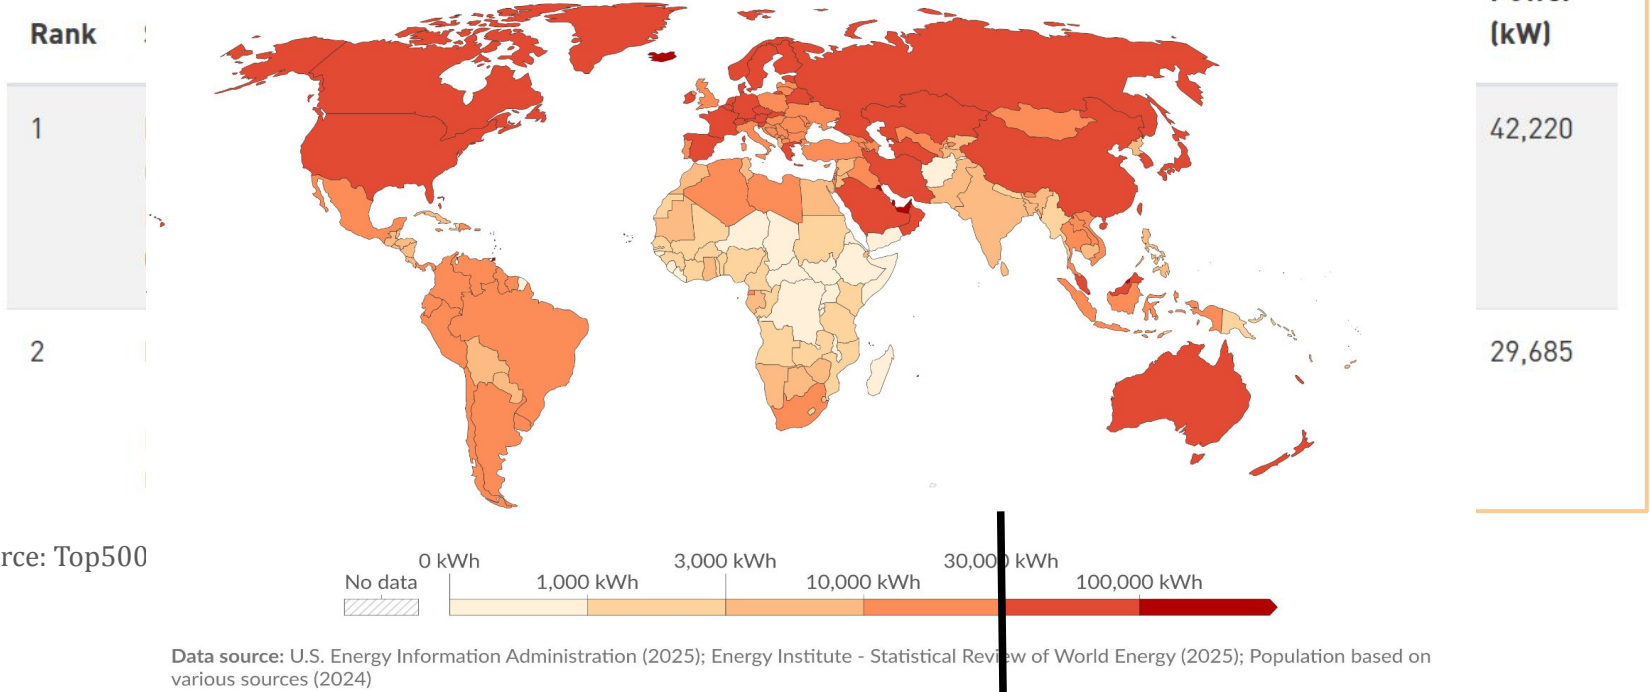
# Energy consumption of supercomputers

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	<b>LineShine</b> - LingKun, LX2 304C 1.55GHz, LingQi, Kylin OS, Shenzhen Cloud Computing Center Co., Ltd. National Supercomputing Centre in Shenzhen (NSCS) China	13,789,440	2,198.40	2,735.82	42,220
2	<b>El Capitan</b> - HPE Cray EX255a, AMD 4th Gen EPYC 24C 1.8GHz, AMD Instinct MI300A, Slingshot-11, TOSS, HPE DOE/NNSA/LLNL United States	11,340,000	1,809.00	2,821.10	29,685

Source: Top500 list, Released June 2026

# Energy consumption of supercomputers

measured in kilowatt-hours per person. Here, energy refers to primary energy using the substitution method.



Source: Top500

Data source: U.S. Energy Information Administration (2025); Energy Institute - Statistical Review of World Energy (2025); Population based on various sources (2024)  
OurWorldinData.org/energy | CC BY

# Pieces of the puzzle

**INFRASTRUCTURE**

# Pieces of the puzzle

**INFRASTRUCTURE**

**HARDWARE**

# Pieces of the puzzle

**INFRASTRUCTURE**

**HARDWARE**

**ALGORITHMS**

# Pieces of the puzzle

**INFRASTRUCTURE**

**HARDWARE**


**ALGORITHMS**

**SOFTWARE**

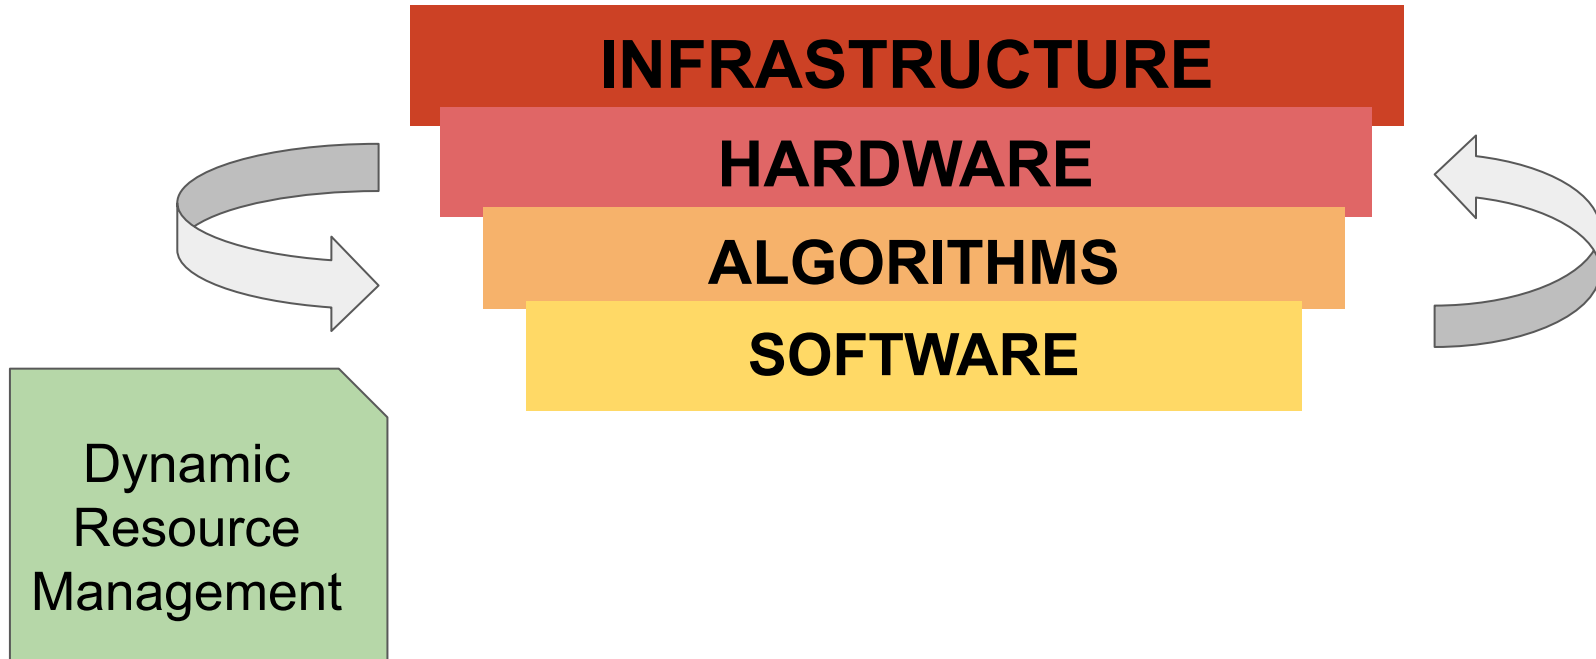
# Pieces of the puzzle



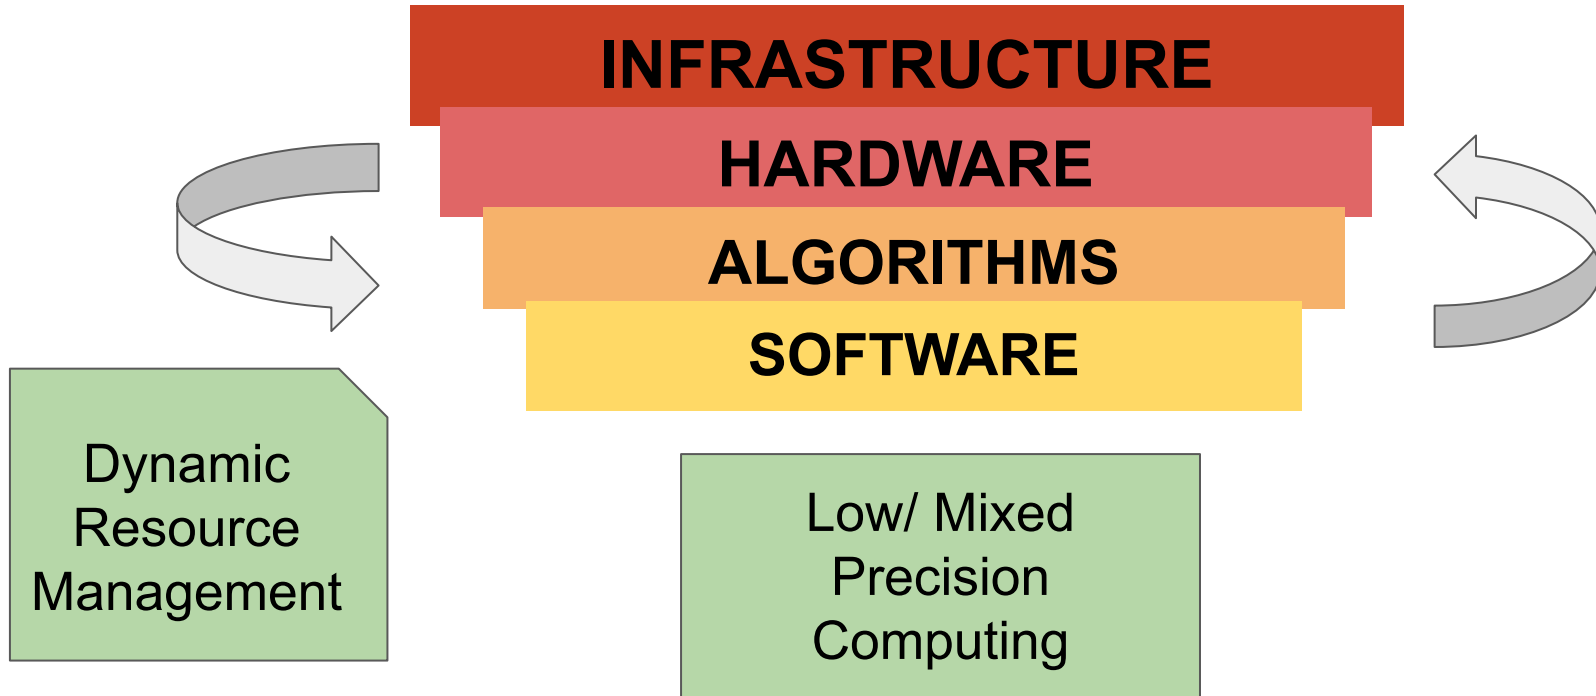
# Pieces of the puzzle




# Pieces of the puzzle



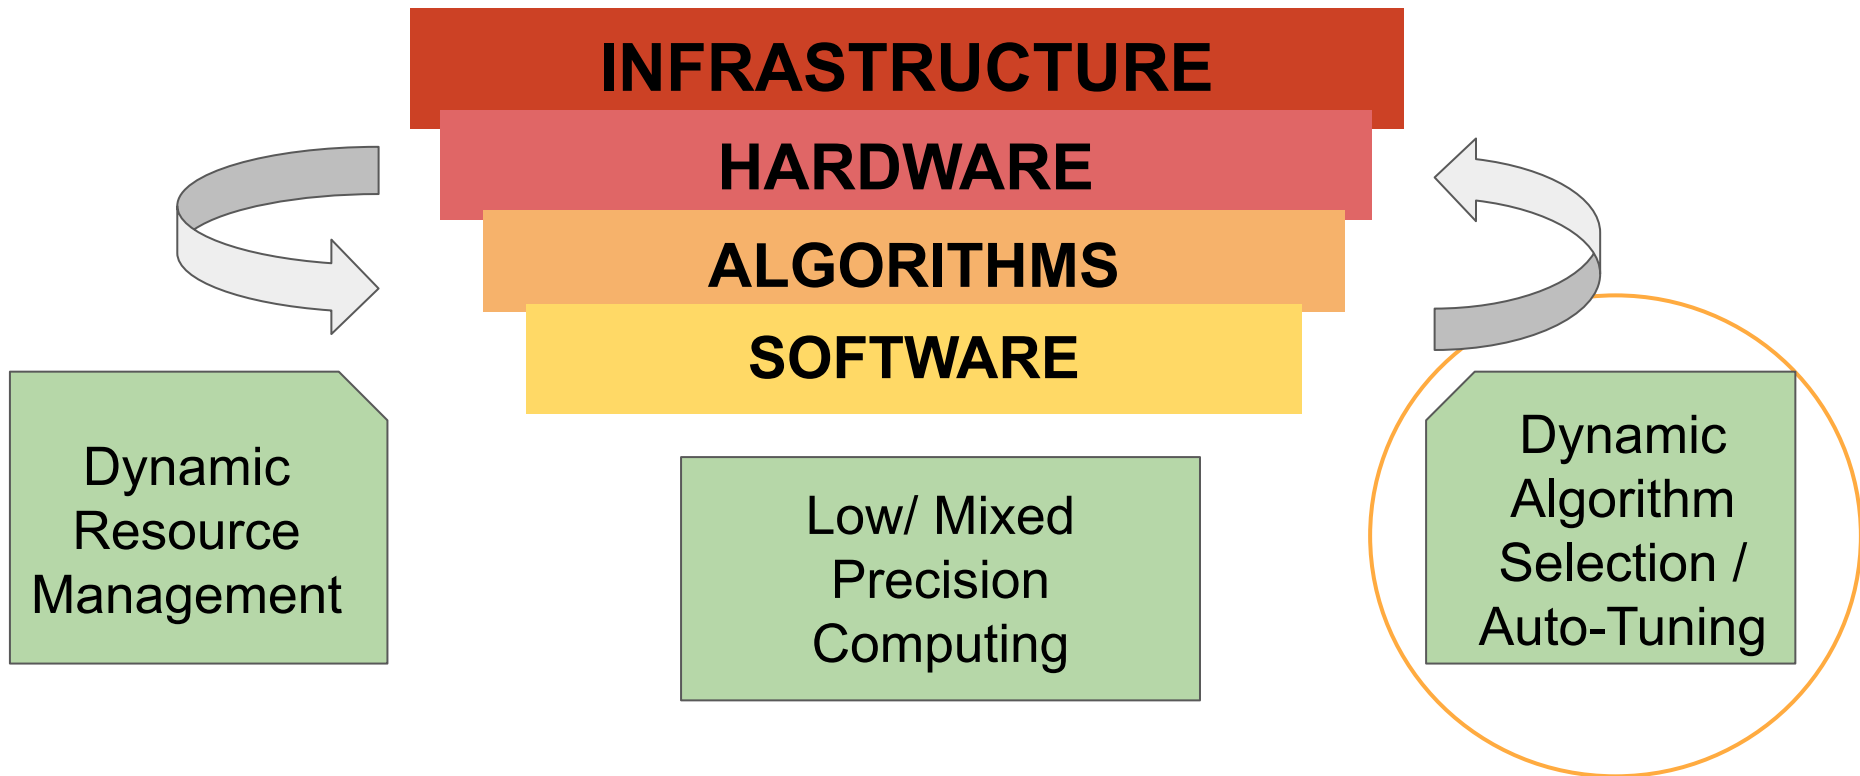
# Pieces of the puzzle



# Pieces of the puzzle



# Pieces of the puzzle



It has already been shown that Dynamic Algorithm Selection for runtime performance can help in reducing overall simulation runtime<sup>1</sup> (and consequently energy)....

<sup>1</sup>Newcome et al., [Algorithm Selection in Short-Range Molecular Dynamics Simulations](#), 2025

It has already been shown that Dynamic Algorithm Selection for runtime performance can help in reducing overall simulation runtime<sup>1</sup> (and consequently energy)....


## Research Question

**Can Algorithm Selection for minimum energy give additional energy savings?**

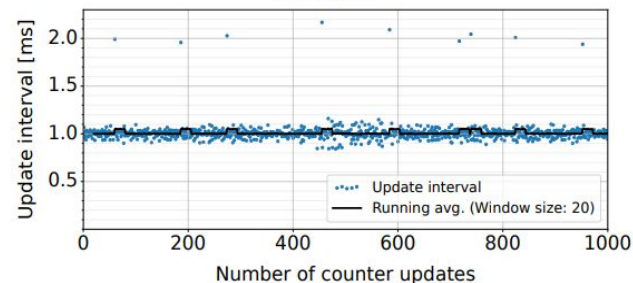
<sup>1</sup>Newcome et al., [Algorithm Selection in Short-Range Molecular Dynamics Simulations](#), 2025

# Energy Measurements on CPUs

- Running Average Power Limit (RAPL)
- Larger error at short intervals (eg. Filtered RAPL)
- RAPL counters update  $\sim 1$  ms
  - Individual kernels can be quite short
  - Needs a mechanism to adjust the interval



**Ice Lake-SP**

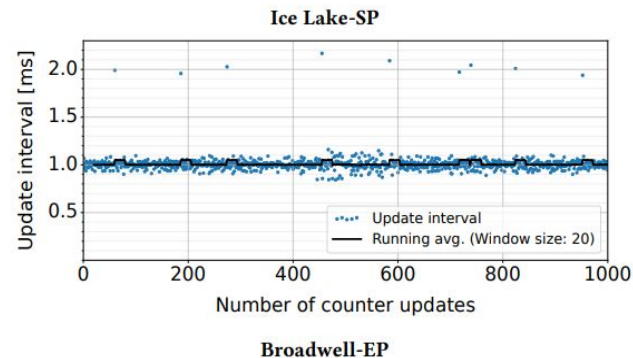



**Broadwell-EP**

Alt et al., An Experimental Setup to Evaluate RAPL Energy Counters for Heterogeneous Memory, 2025

# Energy Measurements on CPUs


- Running Average Power Limit (RAPL)
- Larger error at short intervals (eg. Filtered RAPL)
- RAPL counters update  $\sim 1$  ms
  - Individual kernels can be quite short
  - Needs a mechanism to adjust the interval




Alt et al., An Experimental Setup to Evaluate RAPL Energy Counters for Heterogeneous Memory, 2025

# Energy Measurements on CPUs

- Running Average Power Limit (RAPL)
- Larger error at short intervals (eg. Filtered RAPL)
- RAPL counters update  $\sim 1$  ms
  - Individual kernels can be quite short
  - Needs a mechanism to adjust the interval



Ice Lake-SP



Broadwell-EP

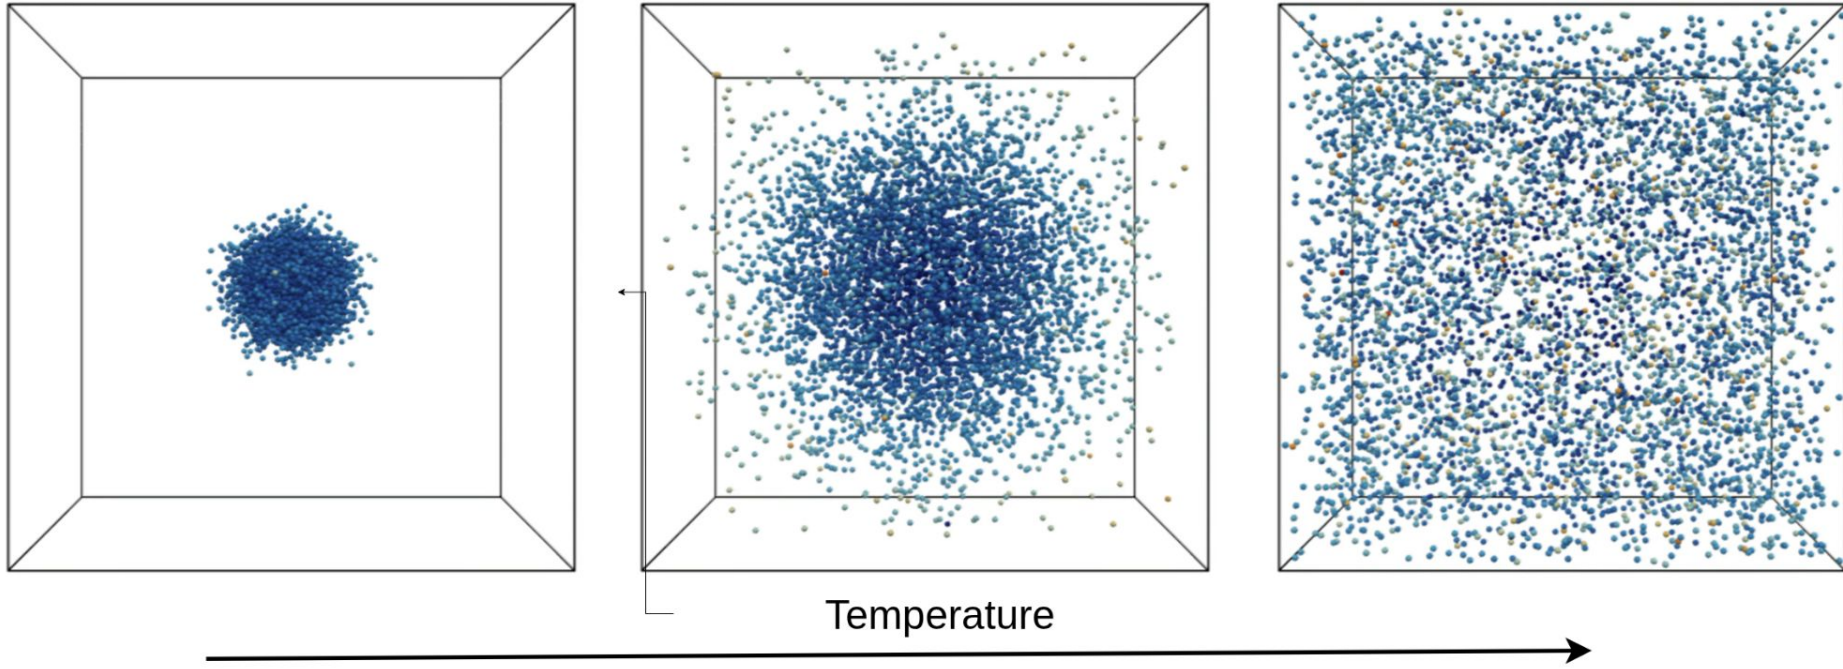
Alt et al., An Experimental Setup to Evaluate RAPL Energy Counters for Heterogeneous Memory, 2025

## Research Question(s)

**Can we replace runtime as a metric with energy and still obtain reliable results?**

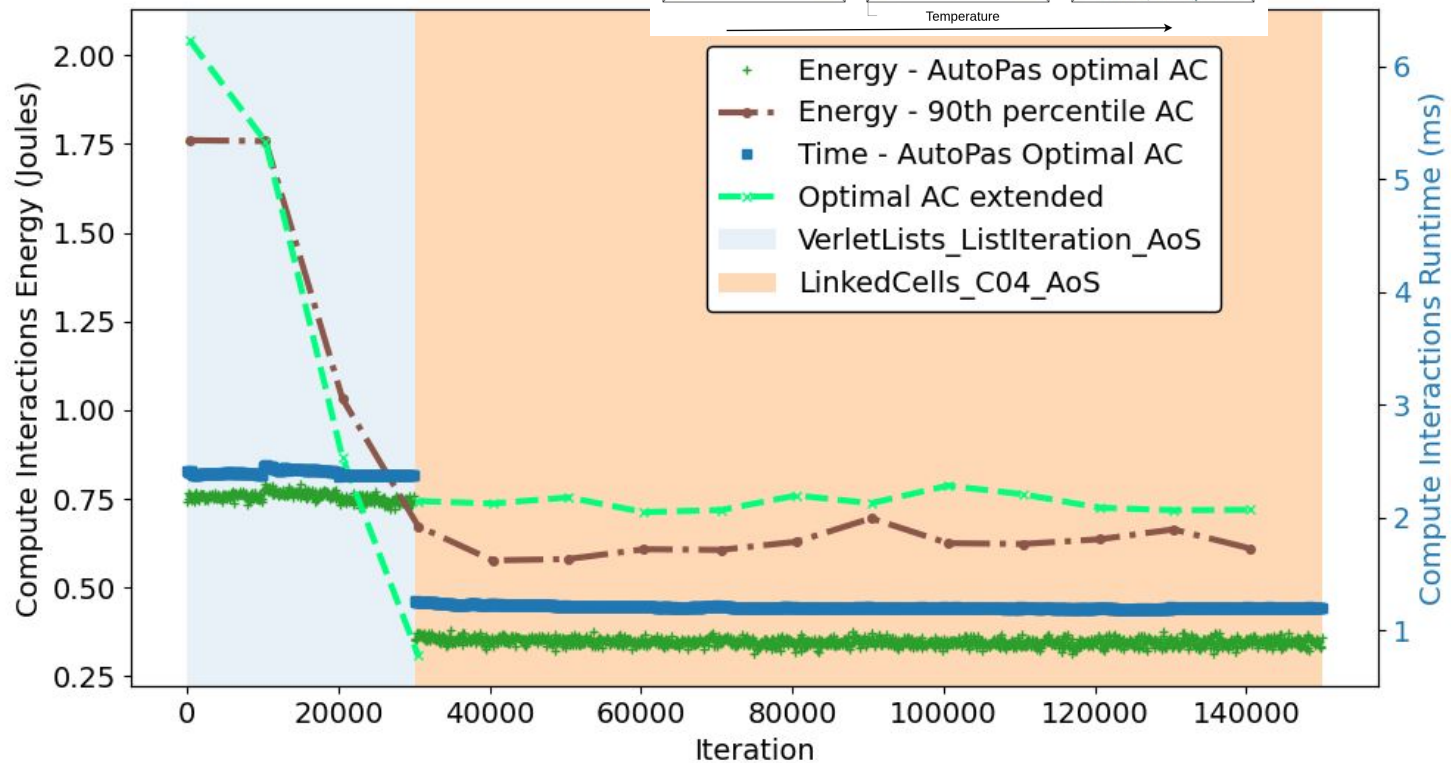

**Can Algorithm Selection for minimum energy give additional energy savings?**

# Heating Sphere

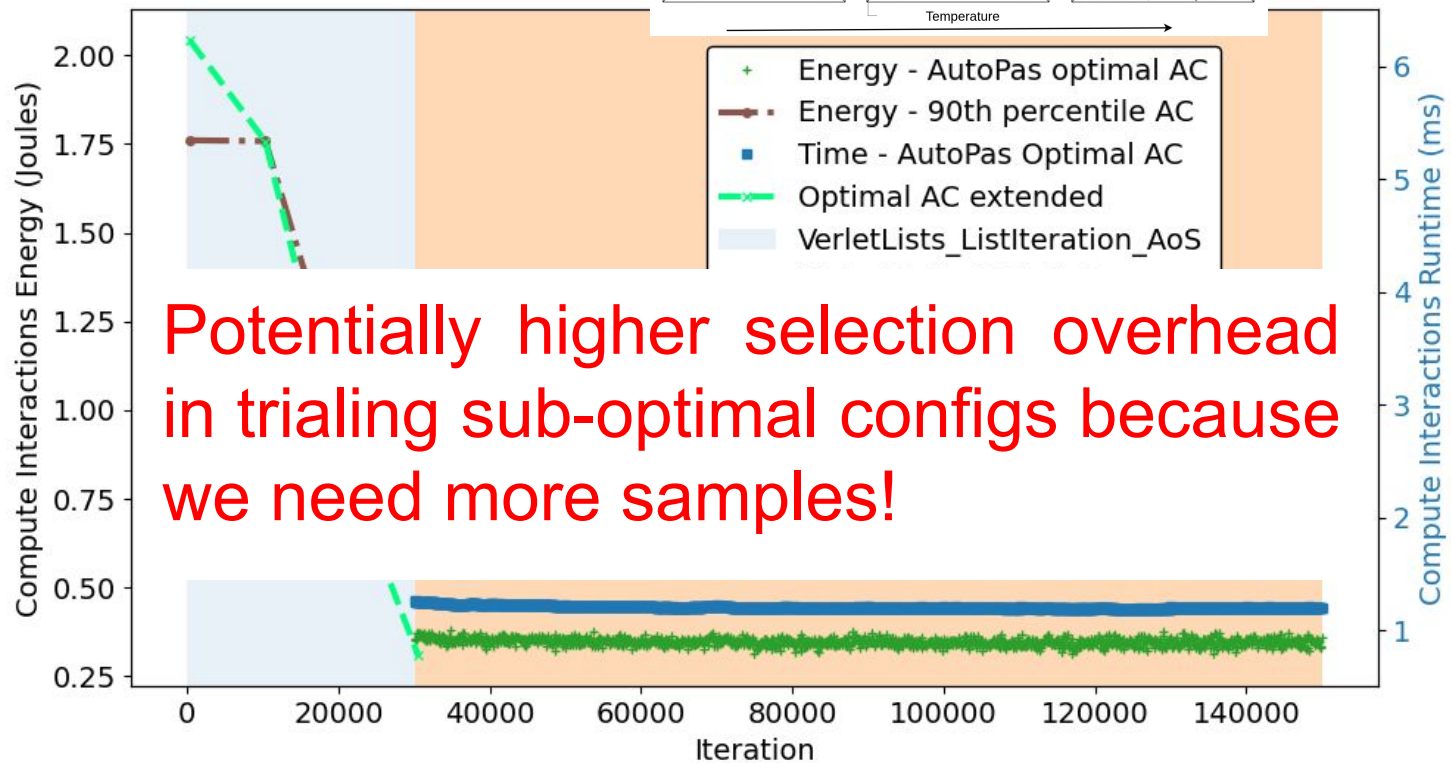



<sup>1</sup>Newcome et al., [Algorithm Selection in Short-Range Molecular Dynamics Simulations](#), 2025


# Heating Sphere



# Heating Sphere



# Dynamic Voltage Frequency Scaling



# Dynamic Voltage Frequency Scaling




Runtime




$AC_1$

Energy




$AC_2$


# Dynamic Voltage Frequency Scaling



Runtime




AC<sub>2</sub>



AC<sub>1</sub>


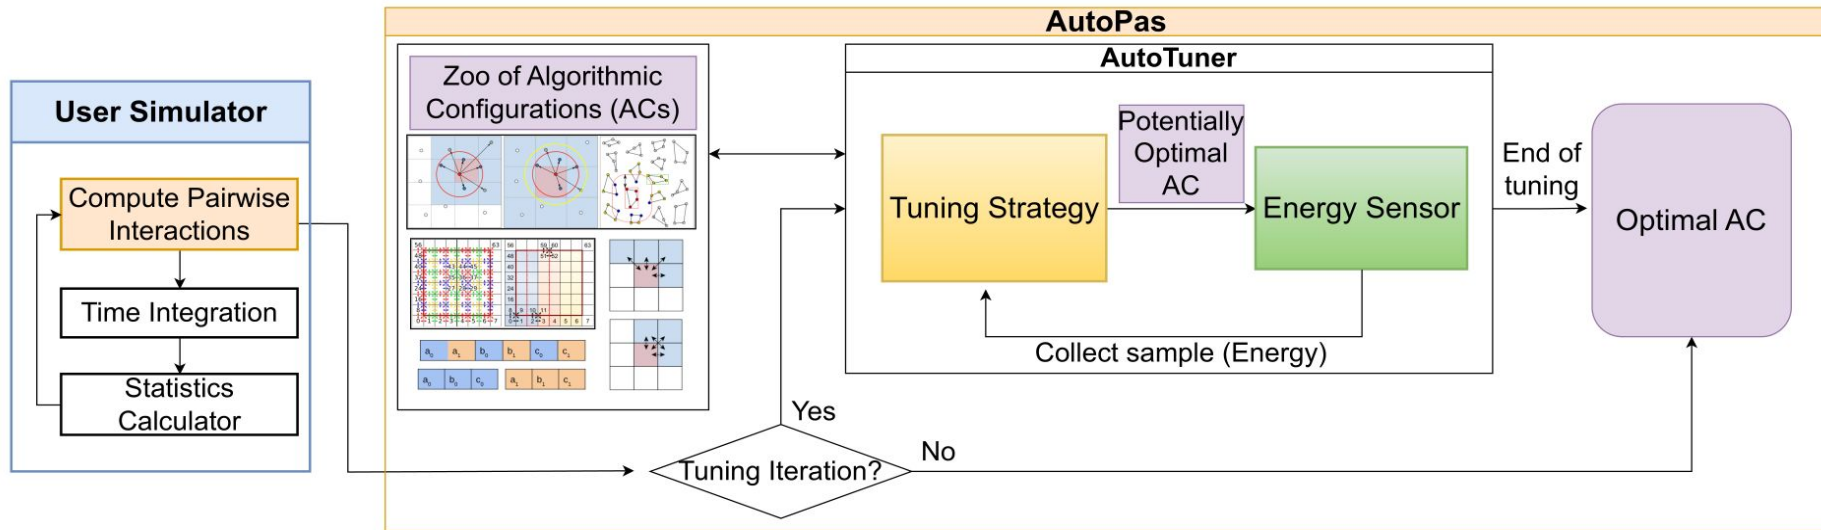
Energy




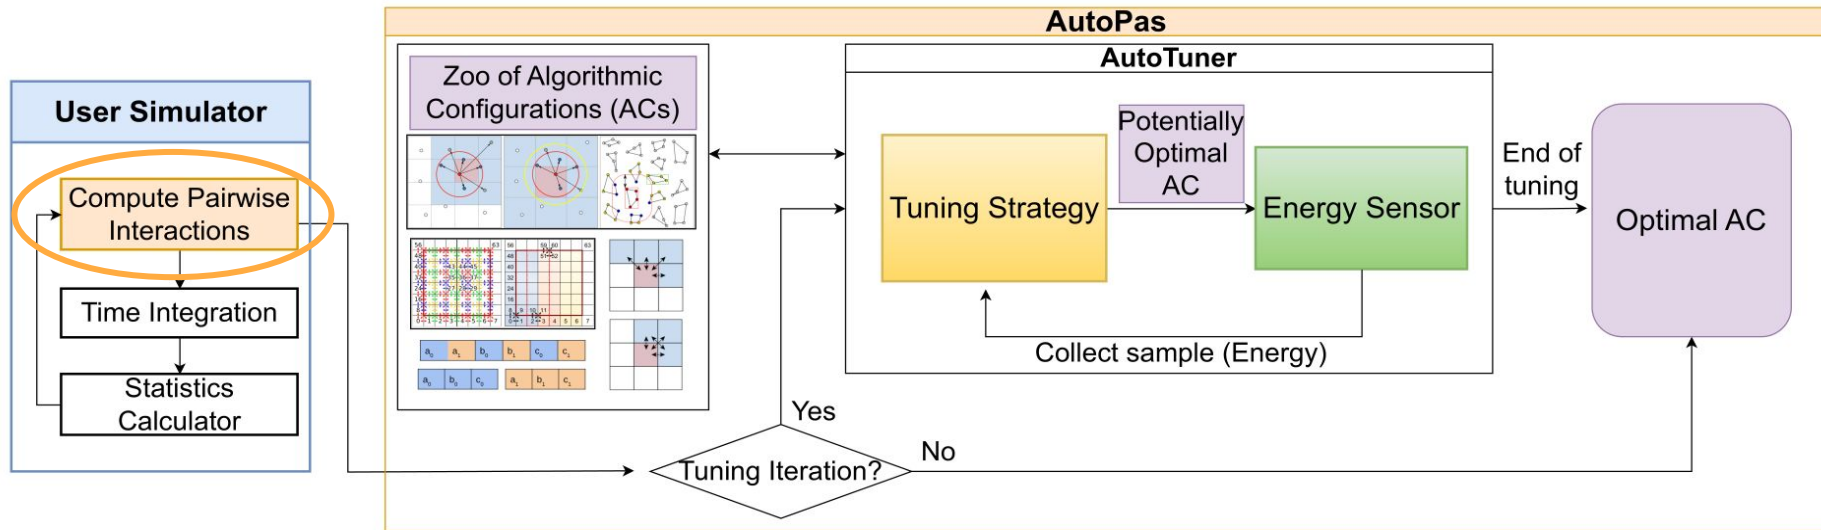
AC<sub>2</sub>

*DVFS can lead to different optimal ACs for runtime or energy!*


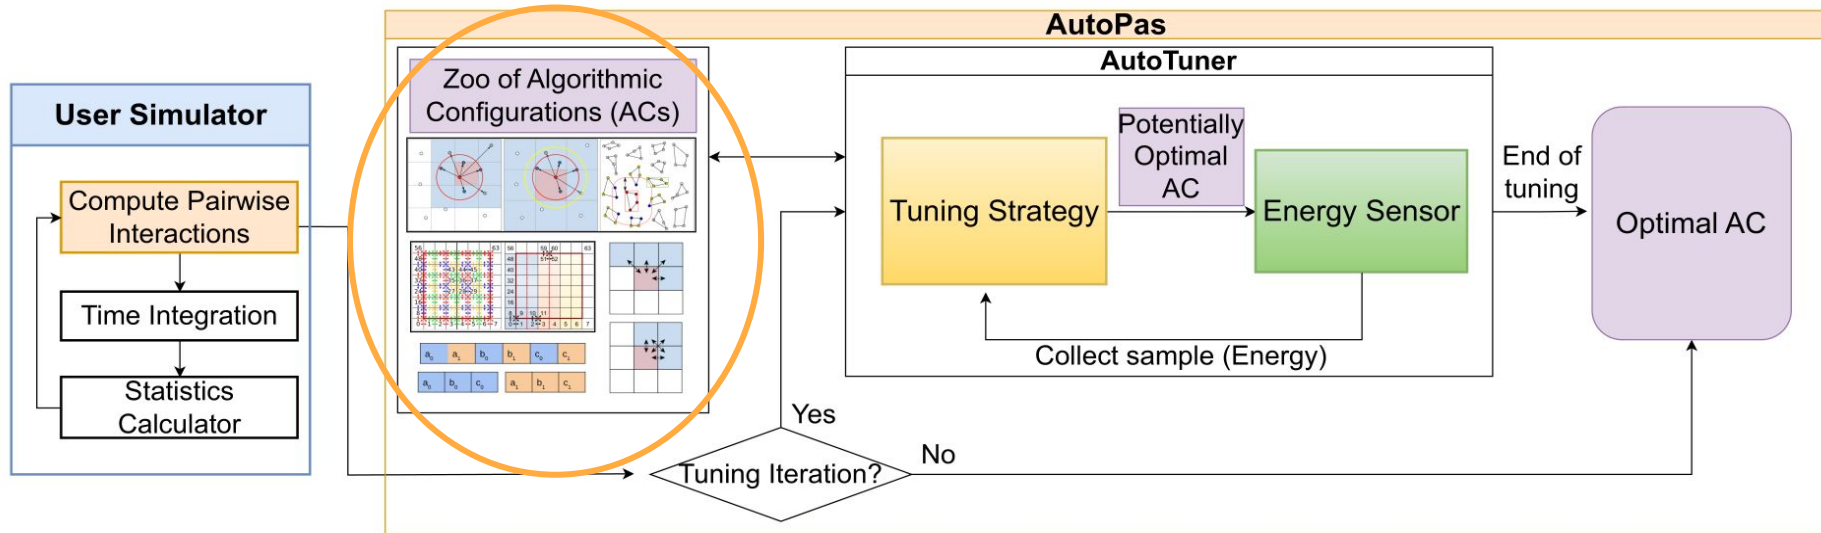
# AutoPas




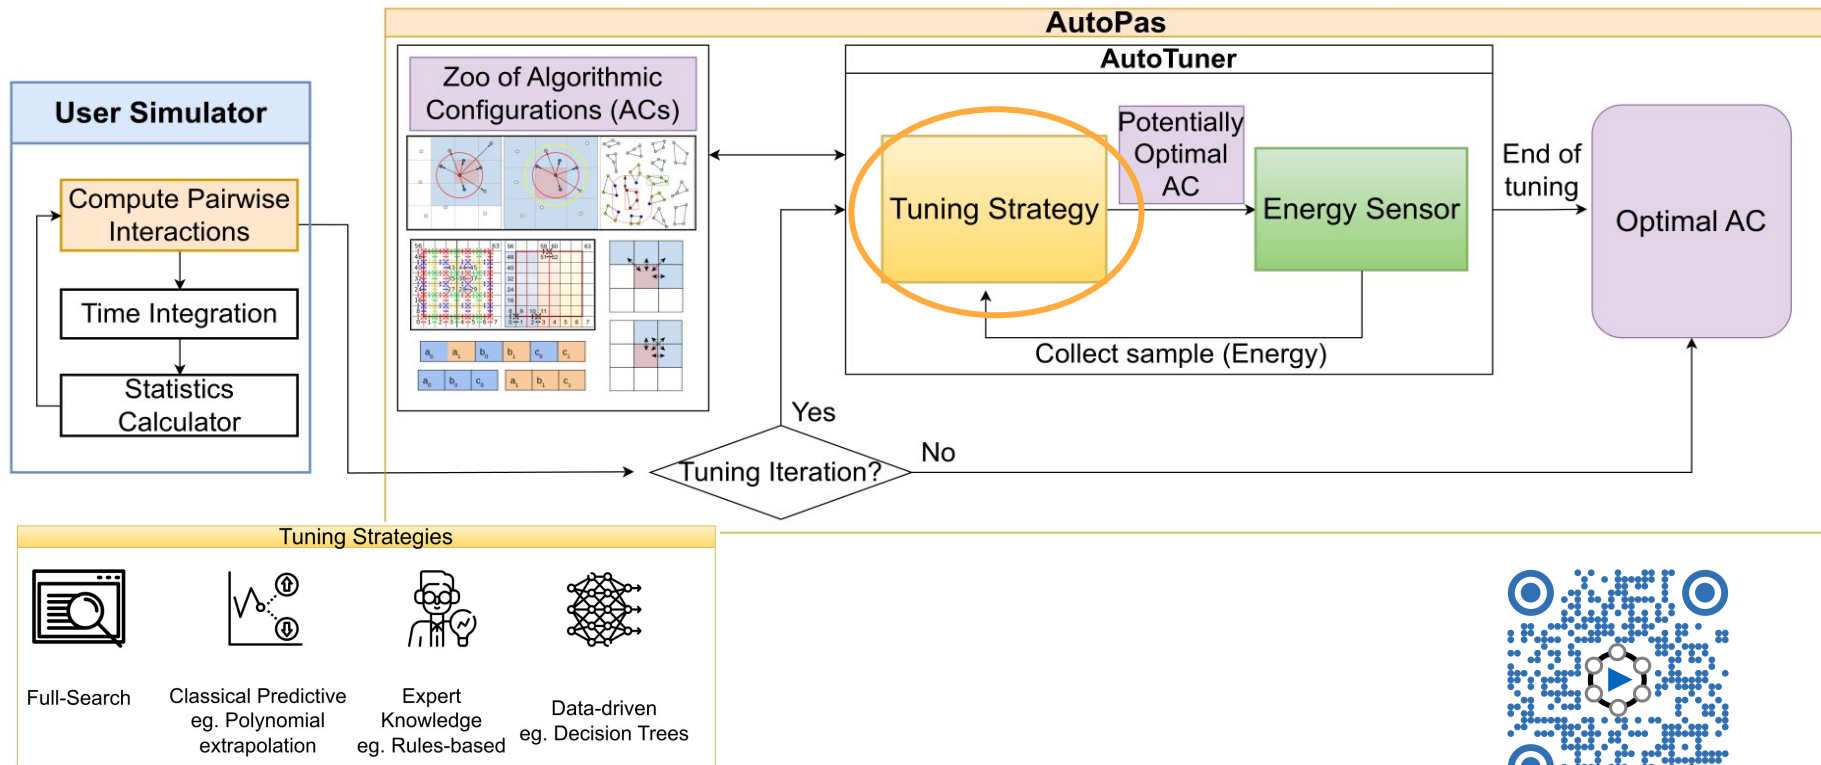
# AutoPas




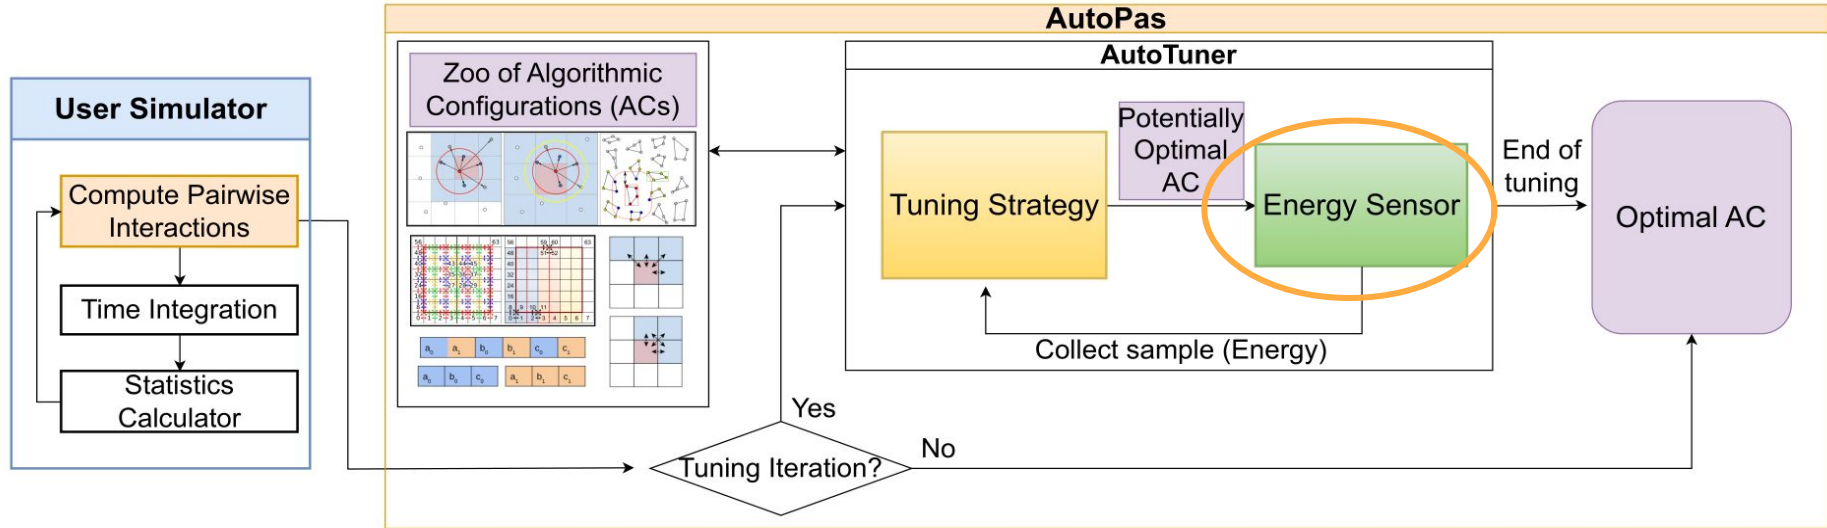
# AutoPas




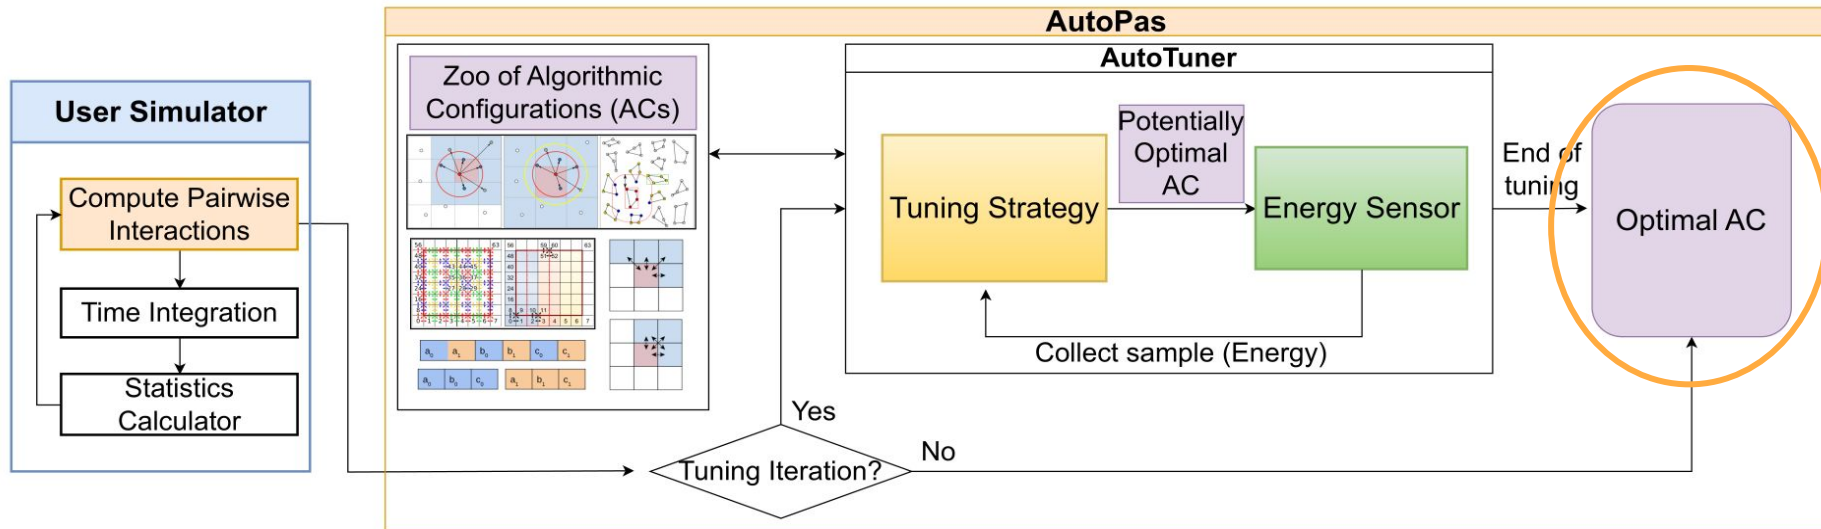
# AutoPas




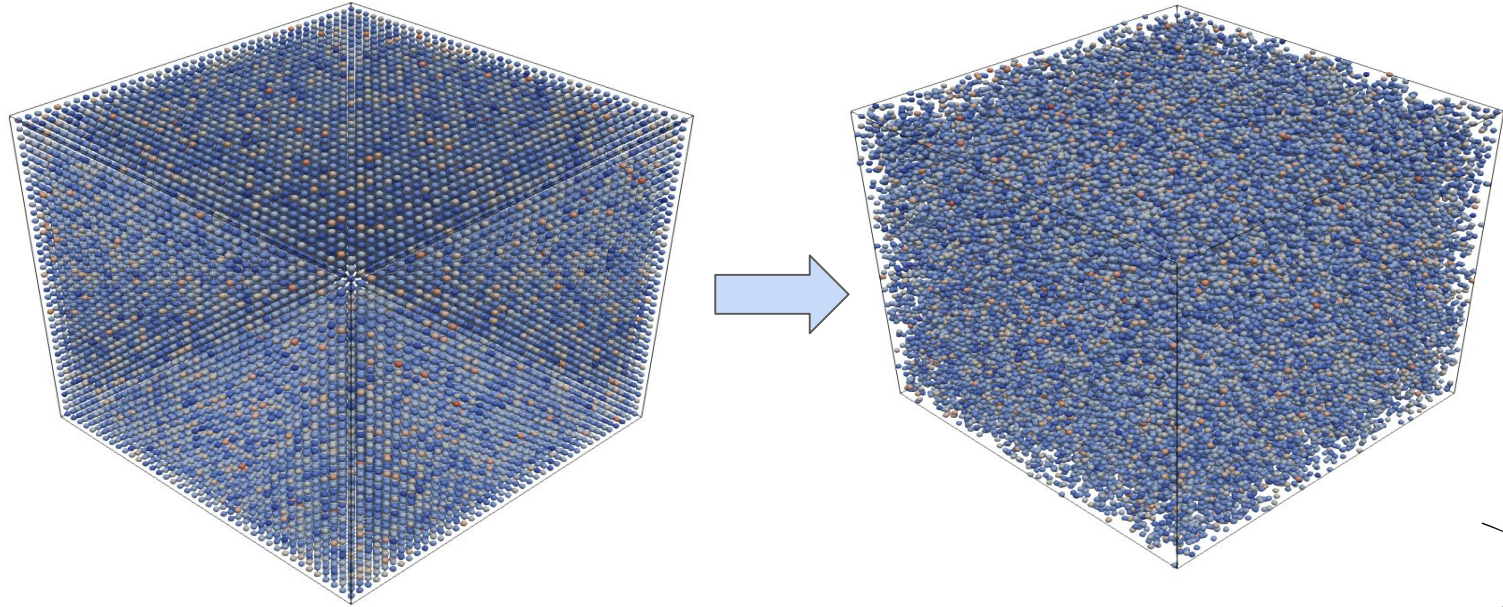
# AutoPas





# AutoPas





# Equilibration Scenario



# Shared-Memory Parallelism





# Shared-Memory Parallelism



C04:

- Dynamic assignment of cell blocks to threads
- Better load balancing

# Shared-Memory Parallelism



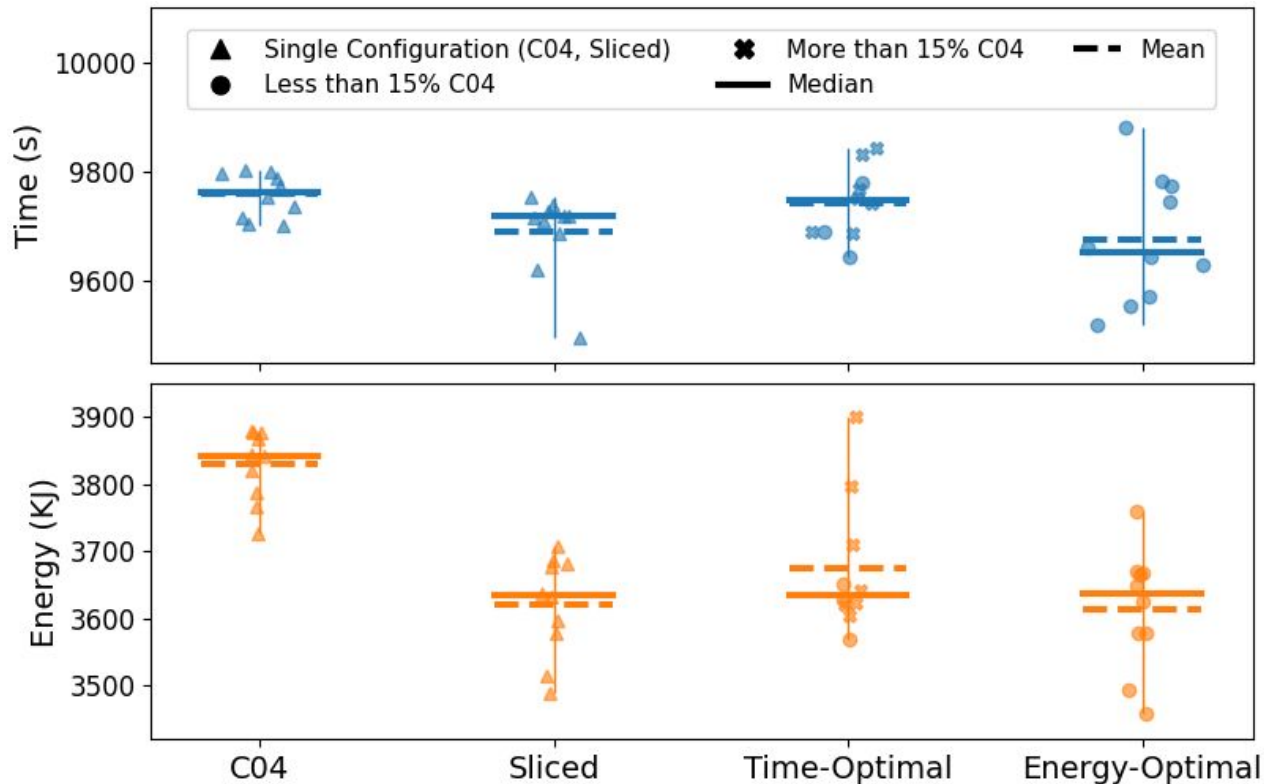
## C04:

- Dynamic assignment of cell blocks to threads
- Better load balancing

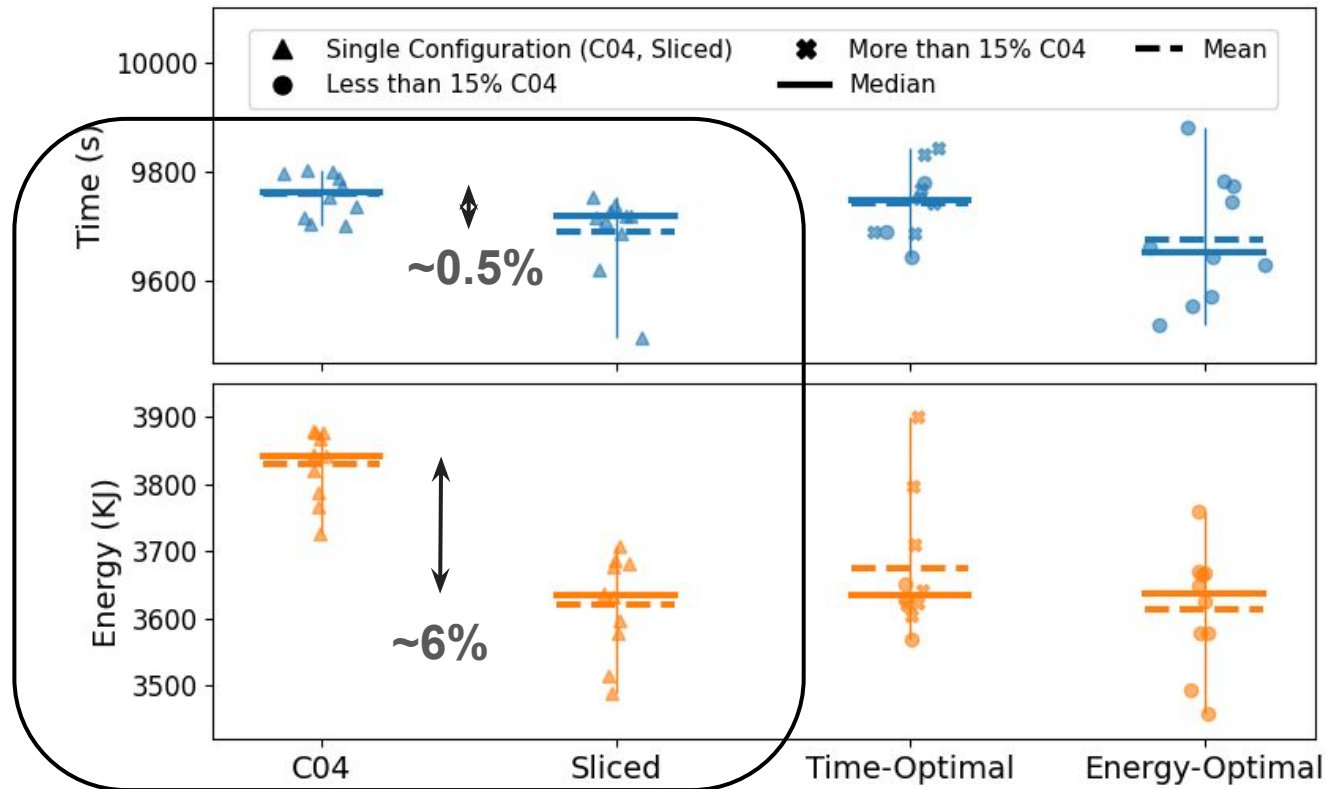
## Sliced:

- Static assignment of cell blocks to threads
- Better memory access pattern

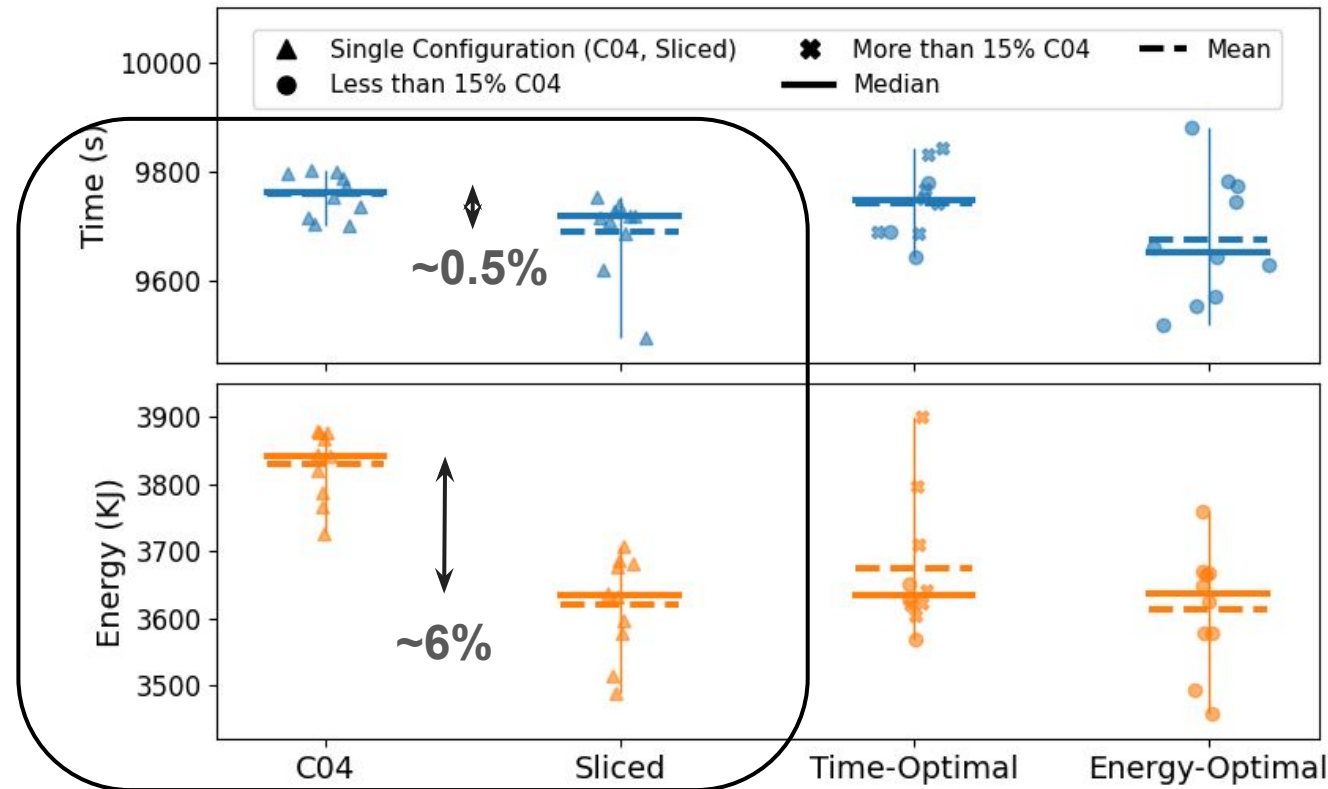
# Results



# Results



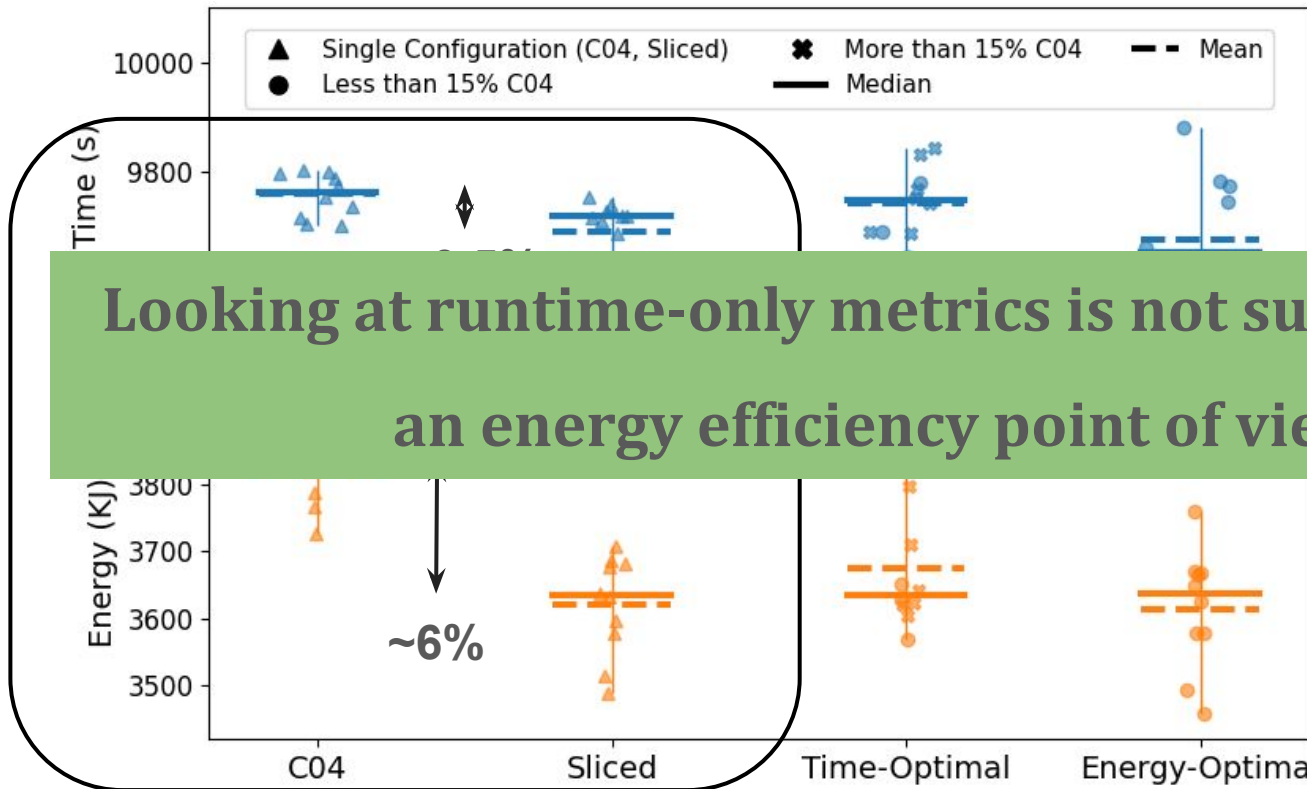
# Results



VTune  
Analysis

*C04 shows higher  
DRAM b/w and  
average CPU  
frequency...  
But lower spin time  
... compared to  
Sliced.*


# Results




VTune  
Analysis

*But lower spin time  
... compared to  
Sliced.*

# Results



# Results




*Sliced always better...  
But C04 is selected sometimes when selecting for runtime... leading to extra energy consumption.*

# Results



**While selecting for time works well most of the time, selecting for energy is more robust metric for energy efficiency.**



*selecting for runtime... leading to extra energy consumption.*


# Summary

- Algorithm Selection can provide improvement in runtime and energy performance.
- Despite challenges in fine-grained energy measurement, Algorithm Selection for minimum energy can find optimal ACs.
- In most cases, selecting or tuning for runtime is sufficient.
- However, looking at algorithm selection strategies targeting minimum energy consumption can in some cases reveal additional savings.

Thank you!  
Any questions?

# Backup Slides



# Short Range Particle Simulations



# Things to keep in mind


Syncing energy measurement with RAPL counter update

→ but overhead for repeated measurements.




Haehnel et al., Measuring Energy Consumption for Short Code Paths Using RAPL, 2012

# The Zoo




## Other parameters:

- Stencils, c04, c08, c18...
- AoS vs SoA,
- Cell-size-factor,
- Newton 3 opt., etc.



Stencils





Coloring

Slicing

# Simulators that use AutoPas

- Ls1-MarDyn<sup>1</sup>
- LADDS<sup>3</sup> - Large-Scale Deterministic Debris Simulation (TUM-ESA collab)
- MD-Flexible, our inhouse Molecular Dynamics Simulator
- LAMMPS with AutoPas<sup>2</sup> (experimental)

**ls1**  
Mardyn



LADDS

<sup>1</sup><https://github.com/ls1mardyn/ls1-mardyn>

<sup>2</sup><https://github.com/AutoPas/lammps-autopas>

<sup>3</sup><https://github.com/esa/LADDS>

# AutoPas features in ls1-MarDyn

- ls1-MarDyn:
  - Particles/Molecules classes, Linked Cells based container, used in chemical engineering
  - MPI + OpenMP hybrid parallelization
  - 20 trillion particle MD simulation on entire set of 9216 nodes of SuperMUC (Phase 1)<sup>1</sup>
- ls1-MarDyn modified to make use of AutoPas features including tuning for minimum energy

**ls1**  
Mardyn



<sup>1</sup>Tchipev et al., [TweTriS: Twenty trillion-atom simulation](#), 2019  
Manish Kumar Mishra M.Sc. | EESP Workshop | ISC High Performance 2026

# Frequency Scaling

$$E = k_1 f^3 t$$

Compute bound:

$$t \approx \frac{1}{f}$$

$$E \approx f^2$$


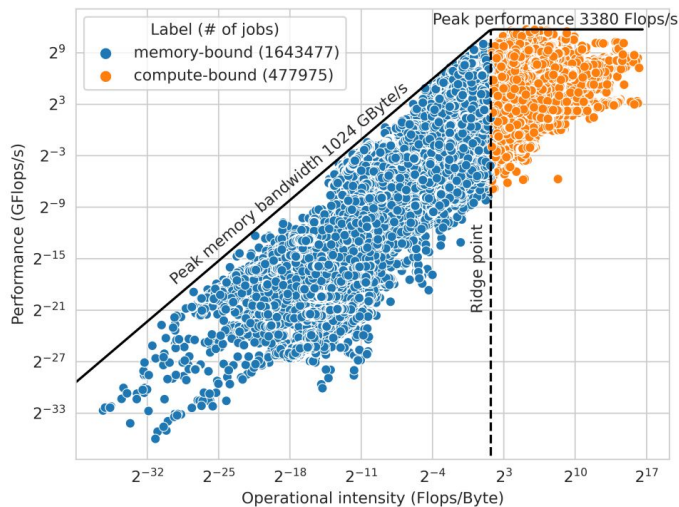
Memory bound:

$$t > \frac{1}{f}$$

$$t \approx \frac{1}{f_r} > \frac{1}{f}$$

$$\Delta f = f - f_r$$

# Frequency Scaling in practice



Antici et al., MCBound: An Online Framework to Characterize and Classify Memory/Compute-bound HPC Jobs, 2024



